

Claims

What is claimed is:

1. A method for providing Web services comprising the steps of:
receiving a request from a Web client process wherein said request includes customer ID information;
spawning a program element operable on a computing node to process said request;
5 associating said customer ID information with the spawned program; and
allocating computing resources of said computing node to the spawned program element in accordance with said customer ID information associated with said request.
2. The method of claim 1 wherein the step of allocating comprises the steps of:
allocating a minimum level of resources to the spawned program element in accordance with said customer ID information.
3. The method of claim 1 wherein the step of allocating comprises the steps of:
allocating a maximum level of resources to the spawned program element in accordance with said customer ID information.
4. The method of claim 1 wherein said customer ID information is encoded in a process name of each said spawned program element.
5. The method of claim 1 wherein said computing resources includes processor time utilization.
6. The method of claim 1 wherein said computing resources includes main memory utilization.

7. The method of claim 1 wherein said computing resources includes secondary storage bandwidth utilization.

8. A system for delivery of services in a client/server distributed environment comprising:

a server computing node;

a server process operable on said server computing node for processing requests from a plurality of client processes coupled to said server computing node;

a plurality of server child processes operable on said server computing node and spawned by said server process to process said requests from said plurality of client processes wherein each child process of said plurality of server child processes is associated with customer ID information; and

a process resource manager operable on said server computing node to control allocation of resources of said server computing node among said plurality of server child processes wherein said resource manager is operable to control allocation of said resources in accordance with said customer ID information associated with said each child process.

9. The system of claim 8

wherein said server process is a Web server process, and

wherein said each child process is a cgi-bin process.

10. The system of claim 8 wherein said process resource monitor further comprises:

a CPU time resource monitor element for allocating CPU time to said each child process;

a secondary storage bandwidth resource monitor for allocating secondary storage bandwidth to said each child process; and

a main memory resource monitor for allocating main memory to said each child process.

17. The system of claim 11 wherein said means for allocating includes means for allocating secondary storage bandwidth utilization.

